



Geospatial Optimization of Siting Large-Scale Solar Projects

Jordan Macknick and Ted Quinby
National Renewable Energy Laboratory

Emmet Caulfield and Margot Gerritsen
Stanford University

Jay Diffendorfer and Seth Haines
U.S. Geological Survey

The Joint Institute for Strategic Energy Analysis is operated by the Alliance for Sustainable Energy, LLC, on behalf of the U.S. Department of Energy's National Renewable Energy Laboratory, the University of Colorado-Boulder, the Colorado School of Mines, the Colorado State University, the Massachusetts Institute of Technology, and Stanford University.

Technical Report
NREL/TP-6A50-61375
March 2014

Contract No. DE-AC36-08GO28308

Geospatial Optimization of Siting Large-Scale Solar Projects

Jordan Macknick and Ted Quinby
National Renewable Energy Laboratory

Emmet Caulfield and Margot Gerritsen
Stanford University

Jay Diffendorfer and Seth Haines
U.S. Geological Survey

Prepared under Task No. 6A50.1010

The Joint Institute for Strategic Energy Analysis is operated by the Alliance for Sustainable Energy, LLC, on behalf of the U.S. Department of Energy's National Renewable Energy Laboratory, the University of Colorado-Boulder, the Colorado School of Mines, the Colorado State University, the Massachusetts Institute of Technology, and Stanford University.

JISEA® and all JISEA-based marks are trademarks or registered trademarks of the Alliance for Sustainable Energy, LLC.

NOTICE

This report was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or any agency thereof.

Available electronically at <http://www.osti.gov/bridge>

Available for a processing fee to U.S. Department of Energy
and its contractors, in paper, from:

U.S. Department of Energy
Office of Scientific and Technical Information

P.O. Box 62
Oak Ridge, TN 37831-0062
phone: 865.576.8401
fax: 865.576.5728
email: <mailto:reports@adonis.osti.gov>

Available for sale to the public, in paper, from:

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
phone: 800.553.6847
fax: 703.605.6900
email: orders@ntis.fedworld.gov
online ordering: <http://www.ntis.gov/help/ordermethods.aspx>

Cover Photos: (left to right) PIX 04135, iStock 22779761, PIX 16933, PIX 15648, PIX 08466, PIX 21205

Printed on paper containing at least 50% wastepaper, including 10% post consumer waste.

Acknowledgments

The authors would like to thank the Joint Institute of Strategic Energy Analysis (JISEA) and its Innovative Research and Analysis Award Program (IRAAP) for providing the funding for this work. The authors would also like to extend appreciation for comments, graphics, and editing support to David Hester of the U.S. Geological Survey, as well as Laura Vimmerstedt, Nate Blair, Doug Arent, Patricia Statwick, Scott Gossett, Mary Lukkonen, and Kendra Palmer of the National Renewable Energy Laboratory. Any use of trade, product, or firm names is for descriptive purposes only and does not imply endorsement by the U.S. Government.

Abstract

Recent policy and economic conditions have encouraged a renewed interest in developing large-scale solar projects in the U.S. Southwest. However, siting some large-scale solar projects, such as concentrating solar power (CSP), is complex. In addition to the quality of the solar resource, solar developers must take into consideration many environmental, social, and economic factors when evaluating a potential site. This report describes a proof-of-concept, Web-based Geographical Information Systems (GIS) tool that evaluates multiple user-defined criteria in an optimization algorithm to inform discussions and decisions regarding the locations of utility-scale solar projects. Existing siting recommendations for large-scale solar projects from governmental and non-governmental organizations are not consistent with each other, are often not transparent in methods, and do not take into consideration the differing priorities of stakeholders. The siting assistance GIS tool we have developed improves upon the existing siting guidelines by being user-driven, transparent, interactive, capable of incorporating multiple criteria, and flexible. This work provides the foundation for a dynamic siting assistance tool that can greatly facilitate siting decisions among multiple stakeholders.

Table of Contents

1	Introduction.....	1
2	CSP Siting	2
	2.1 CSP Siting Challenges	2
	2.2 Previous Approaches to Develop Siting Criteria	2
3	Results.....	6
	3.1 Overview of Siting Assistance Tool	6
	3.2 How a User Interacts With the Tool	6
4	Methods.....	9
	4.1 System Design	9
	4.2 Site Selection Algorithm.....	9
5	Discussion.....	14
	5.1 Proof-of-Concept Tool.....	14
	5.2 Future Directions and Recommendations	14
6	Conclusions	16
	References.....	17
	Appendix: System Technical Description	18
	System Structure.....	18
	Operation	19
	Google Earth Development Client	20
	OpenCarto User Interface.....	20

1 Introduction

Current demographic trends indicate a surge in population and accompanying energy demands in the arid U.S. Southwest (MacDonald 2010). New sources of electricity must meet this increasing demand and replace the retiring fleet of existing electricity generating units in the region. The Energy Improvement and Extension Act (2008) and the American Recovery and Reinvestment Act (2009) created incentives and a renewed interest in large-scale solar projects with both developers and investors. The acts provided investment tax credits, federal grants, possibilities for federal loan guarantees, and accelerated depreciation. Concentrating solar power (CSP) is one type of large-scale solar technology with the potential to generate considerable amounts of renewable electricity in the region (Mehos et al. 2009;; SEIA 2014; DOE, 2012).

Decisions regarding locations of CSP facilities are complex and involve many factors that often conflict with one another. Developers must consider a variety of factors, including solar resources, proximity to existing infrastructure, environmental regulations and impacts, and cultural heritage sites. In turn, local stakeholders, regional environmental groups, and government agencies also consider these factors when evaluating proposed CSP developments, but may prioritize these factors differently. As a result, some locations selected by developers may not be deemed suitable by other groups and vice versa.

In order to facilitate CSP siting, governmental agencies and non-governmental organizations (NGOs) have determined zones where CSP development would be most suitable. Despite having similar aims, these agencies and groups have utilized different metrics and different weighting schemes to come to their conclusions, resulting in zones that do not overlap, and mixed signals being sent to CSP developers over high priority solar sites. These differing priorities led to many criticisms by local, regional, and national NGOs in the zones developed by the Bureau of Land Management (BLM), eventually resulting in some of the zones being modified (BLM/DOE 2012). Furthermore, because fundamental assumptions and methods are not fully transparent, it is difficult to understand how or why different efforts produce different results. An essential component of decision making is effective dialogue among stakeholders to ensure all viewpoints are considered (Clark 2002). However, in the case of large-scale solar siting procedures, there is a lack of common metrics and tools that can be used among all stakeholders.

There is a need for a common framework whereby ideas from different stakeholders can be considered. We have developed a proof-of-concept siting assistance tool that transparently weighs various factors that influence decisions of CSP plant locations and allows users to alter the weighting mechanisms according to their stated preferences. We provide the rationale and foundation for a tool utilizing existing Geographic Information System (GIS) data that can successfully address a subset of the factors considered in CSP siting decisions. More work is required to finalize the tool and expand its scope beyond CSP siting concerns to include other energy and mineral resource development activities.

2 CSP Siting

2.1 CSP Siting Challenges

Identifying suitable sites for CSP deployment requires fulfilling multiple objectives. A threshold of solar direct normal insolation (DNI) of around 6 kWh/m²/day is generally required to produce enough electricity economically. In the United States, this quality of solar insolation is generally only found in the Southwest (Mehos et al. 2009). In addition, large areas of contiguous land (5 to 13 acres for every megawatt of capacity) are required to account for the size of typical proposed CSP facilities (BLM/DOE 2010). This land must also be relatively flat. The degree of slope allowed will vary based on CSP technology, but generally 5% land slope is considered the maximum allowed, while most developers look for slopes less than 2% (BLM/DOE 2010). Protected areas — national parks and preserves, wilderness areas, wildlife refuges, and critical habitats for endangered species — and buffer zones around these areas are excluded from consideration. Developers also consider proximity to existing infrastructure, such as electric transmission lines, roads, and urban areas, as these can significantly impact overall project costs and expected environmental impacts. Existing land uses and land ownership must also be considered, as these factors could influence how difficult it is to obtain leases and permits in certain areas. Other siting considerations include the likelihood of visual impacts on culturally significant viewsheds, the environmentally sensitive habitats in the area, and the physical and legal availability of water resources for plant operations. The effects of development on local communities related to economic development, environmental justice, and available social support must also be considered.

2.2 Previous Approaches to Develop Siting Criteria

Governmental agencies and NGOs have attempted to take these multiple criteria into account to develop appropriate guidelines for CSP siting. Some have taken a more prescriptive approach and have developed maps with explicit zones suitable for development, whereas others have taken an exclusionary approach, listing areas that should be avoided in development.

Mehos et al. (2009) utilized a map overlay method to showcase the amount of available land and solar generation potential for CSP (Figure 1). They developed the theoretical potential for CSP deployment by considering DNI greater than 6 kWh/m²/day, excluding areas with land slopes >1%, legally protected environmental areas, urban areas, and water bodies. This land area corresponds to an installed CSP capacity of about 7,000 gigawatts (GW), or approximately 16,000 gigawatt-hours (GWh) of electricity per year, approximately four times the total national electricity consumption in 2009 (4,119 GWh) (Mehos et al. 2009).

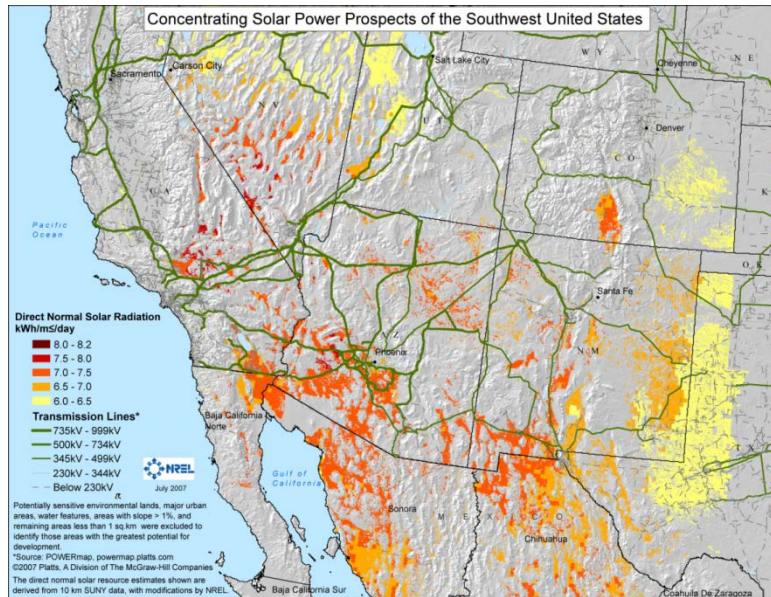


Figure 1. Theoretical land potential for CSP development under solar, land, urban, water body, slope, and legal environmental constraints (Mehos et al. 2009).

The Western Governors' Association (WGA) identified geographic areas called Western Renewable Energy Zones (WREZ) that are suitable for solar, wind, hydropower, and geothermal installations (WGA/DOE 2009). These sites incorporated various metrics, including protected environmental areas, land slope, land ownership, transmission infrastructure, and energy resources (Figure 2).

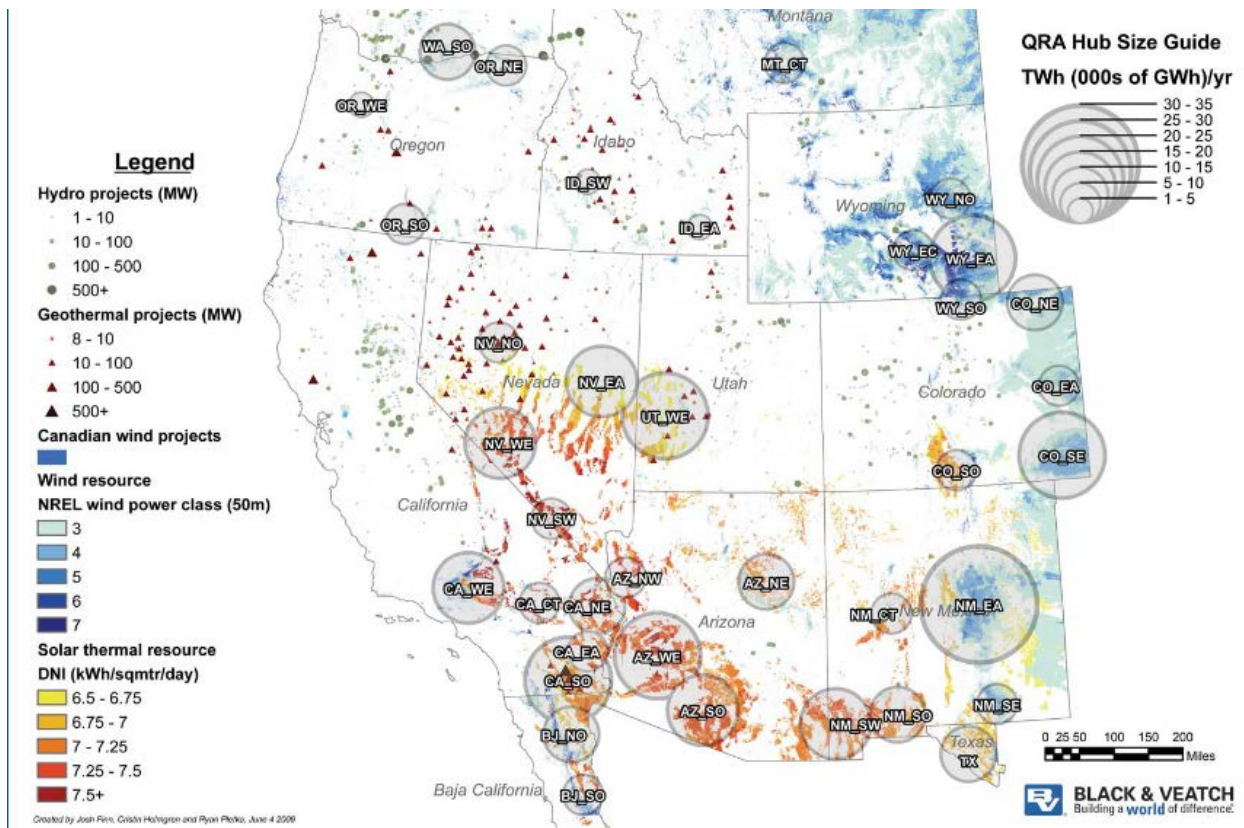


Figure 2. Western Renewable Energy Zones (WREZ) (WGA/DOE 2009). QRA = Qualified Resource Area. TWh = Terawatt-hour.

The Department of the Interior proactively assigned potential BLM lands for solar project development and is refining a nationwide Solar Energy Development Programmatic Environmental Impact Statement (PEIS). The PEIS initially resulted in 24 proposed solar energy zones (SEZ) on BLM land; that number has been reduced to 17 (BLM/DOE 2012) (Figure 3). By 2030, the PEIS calls for 24 GW of utility-scale solar projects on 200,000 acres of BLM land.

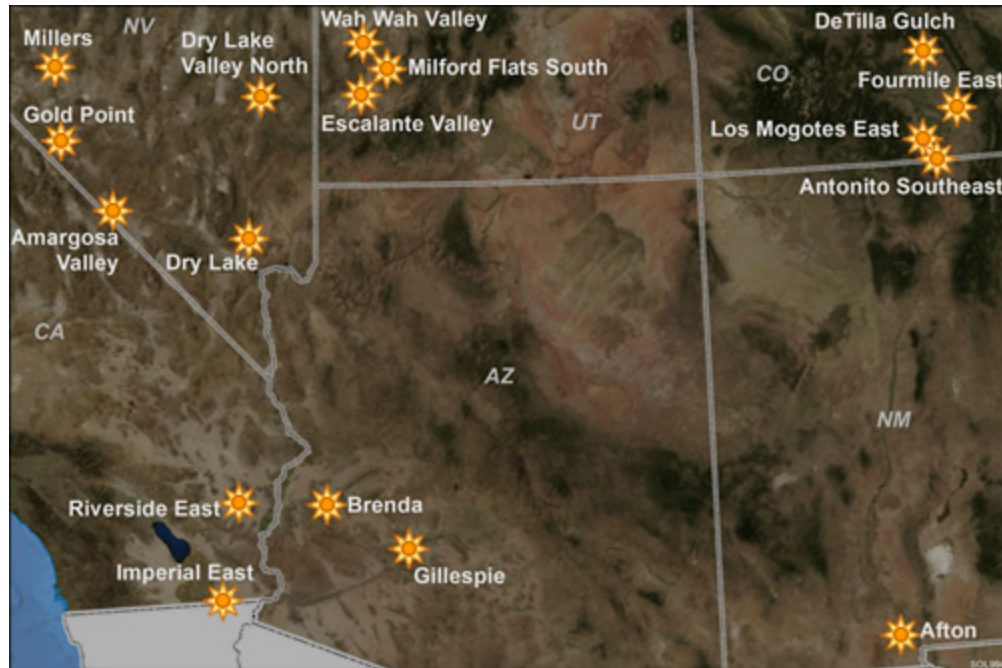


Figure 3. Proposed zones from the PEIS (BLM/DOE 2012).

The California Energy Commission (CEC) released guidelines for siting renewable projects in desert areas in 2009. These guidelines, written for California but applicable to similar ecosystems elsewhere, recommended prioritizing already degraded lands while focusing primarily on identifying and avoiding environmentally sensitive areas (CEC 2009).

In addition to proactive recommendations of solar installation locations, a number of NGOs have responded to these recommendations with alternative maps or guidelines. Each recommendation utilizes a different set of criteria, and each of their prescriptive and prohibitive recommendations differ in important qualitative and spatial ways. The areas excluded tend to differ, as do the degrees of importance they place on certain types of land. For example, some areas, such as wetlands, are treated in three different ways. Some organizations request wetlands development be prohibited (TWC 2009; WGA/DOE 2009), others ask that developers “avoid” these areas (CEC 2009), and still others request that developers “consider” the impacts in these areas (TWS 2009).

Importantly, all of the guidelines described above are static analyses utilizing data in a simple, overlapping method to reveal areas suitable for development. What these studies and guidelines do not provide is a way to compare various potential sites within the allowable area, nor do they allow for alternative valuation systems to account for differences in preferences among stakeholders.

The following section describes the proof-of-concept tool that incorporates these metrics in an optimization algorithm and has a transparent weighting system that incorporates users’ values.

3 Results

3.1 Overview of Siting Assistance Tool

We developed a proof-of-concept for an Internet-based, interactive portal designed to improve the process of siting energy development by incorporating six key elements.

- **Transparency:** Underlying geodata, methods, and assumptions determine the outputs of any decision-making process. Unfortunately, documenting these in sufficient detail to allow perfect replication is extremely difficult. We designed our siting tool to always implement the same algorithm on the same set of data but with user-defined weighting regimes. By recording both the selection of input variables and the resulting outputs, the tool maintains clarity.
- **User-driven:** Differences of opinion in a planning process are difficult to integrate into a final decision. The tool allows every stakeholder to generate siting maps using criteria they specify.
- **Web-based and interactive:** Harnessing web technology allows greater participation in the planning process, simplifies data management, minimizes the potential for differences in methodologies, and facilitates transparency.
- **Multi-criteria optimization, not map overlays:** Spatial overlays of geodata help define and narrow the location choices for siting new development. However, spatial overlays have two shortcomings when used for decision support. First, users quickly fall into a “what if” trap: “What if we overlay these layers instead of those?” This issue quickly becomes an exercise in selecting, weighting, and overlaying maps based on value-laden criteria. Second, decision makers get overwhelmed by a deluge of maps, each reflecting different combinations of data layers, without the underlying reasons for the choice of layers explained or justified. A preferred method is to allow users to first state their preferences or needs as they relate to siting new energy development and then consistently use an optimization algorithm to solve the siting problem.
- **Use existing data:** For the majority of U.S. locations, sufficient geodata exist to support multi-criteria decision making. Although there are some discrepancies in geospatial resolution and data vintage, these do not preclude analysis. We developed a methodology specifically to use existing data to (1) avoid the high costs and time delays associated with data creation and (2) enhance transparency by taking advantage of the quality control and data standards associated with many existing datasets.
- **Flexibility:** As new data layers are developed, they can be incorporated into versions of the tool. This could facilitate further analysis in specific geographic areas where more environmental data layers exist, for example.

3.2 How a User Interacts With the Tool

In the Web-based interface, users can initially select and explore the study area, view the underlying geodata, and move or zoom across the landscape using typical Web interface tools in OpenCarto (see Appendix A). The interface has a drop-down menu called “siting tool” that users select to begin the modeling process.

The siting tool tab opens up a selection menu, allowing the user to specify the criteria for each data layer in the overall siting model. Currently, the tool uses the following layers: solar resource, land slope, transmission lines, existing solar energy development zones, roads, BLM Areas of Critical Environmental Concern (ACEC), critical plant habitat, and critical fauna habitat (Figure 4). Each data layer has an associated cost function based on overlap (e.g., whether it resides in an existing solar development zone) or proximity (e.g., distance to nearest transmission line) to a potential site.

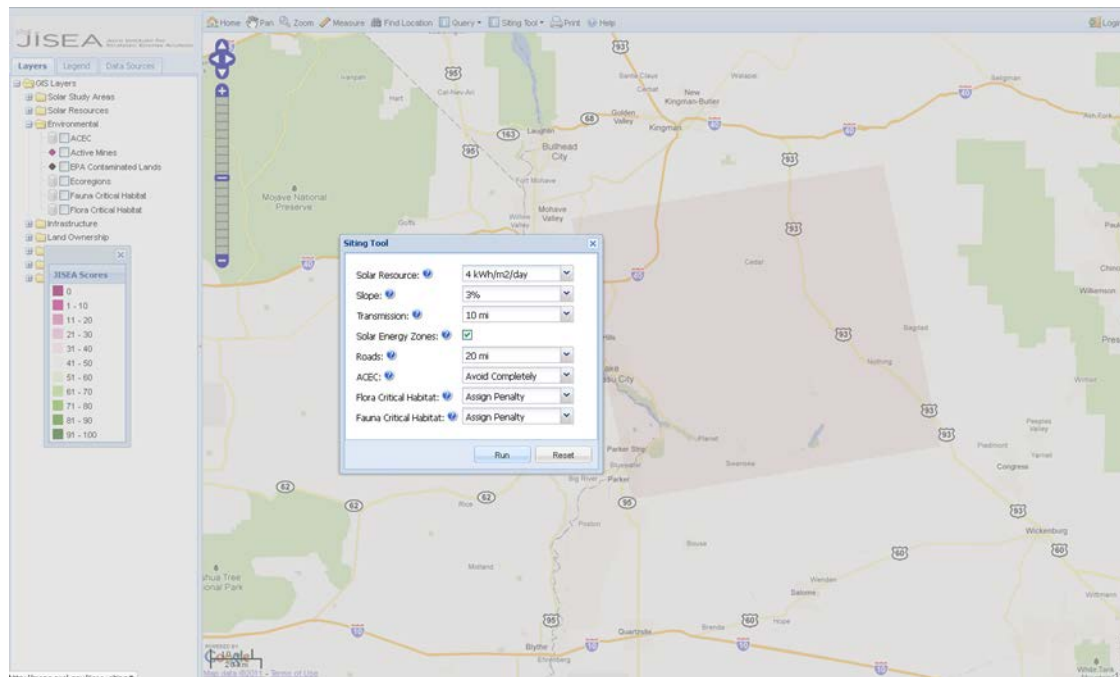


Figure 4. User interface of siting assistance tool. User-defined siting criteria are shown for an illustrative potential development area.

For each of these data layers, a user is presented with a drop-down menu of choices. The choices reflect the type of information in the data layer and allow the user to specify how each layer should be used in the decision-making process. For example, the solar resource layer allows users to specify a threshold amount of solar resource they require for development (in kWh/m²/day). Values on the map above the specified threshold are weighted higher than values below the threshold.

We created weighting regimes for the data layers primarily to reflect different types of mathematical functions, allowing us to test the optimization algorithm. For example, the critical habitat menus include the options of using either a binary weighting regime (i.e., entirely avoid land with critical habitat) or a cost penalty function. The road layer specifies a minimal acceptable distance, beyond which areas are considered unfeasible based on current economic conditions and market demand. Alternatively, the road function could be designed to estimate true costs of new road development. The flexibility of this tool to incorporate binary, cost penalty, and distance-related functions allows multiple types of criteria to be applied to the datasets and incorporated into the analysis.

After selecting criteria preferences, the siting tool outputs a map (with downloadable data) displaying a symbolized “score” for each pixel in a chosen area (Figure 5). The score ranges from 0 to 100 and is based on the user-specified values for each variable. This map can be used in comparisons with existing projects and solar development zones. In addition, location information can be cross-referenced and used on other software programs designed for large-scale solar analysis, such as the Solar Power Prospector (<http://maps.nrel.gov/prospector>) and the System Advisor Model (<http://sam.nrel.gov>).



Figure 5. Example tool output. Green areas indicate locations with higher scores, indicating the best match with user-defined siting criteria.

4 Methods

4.1 System Design

The system includes four main components (Figure 6). Users interact with the system at their computers via a web-based platform. Information comes to the user from two different servers. One delivers maps, images, and the underlying geodata using the National Renewable Energy Laboratory's (NREL) OpenCarto infrastructure. The analysis server performs the site selection based on (1) information specified by the user and (2) the geodata existing on the map server. It then returns the information to the user via a Web-based analysis service (WAS). The analysis program is implemented in C++. A final component is invisible to the user but allows the developers access to the system to alter code, update geodata, and perform general maintenance. Appendix A provides a detailed description of the information technology used in the system.

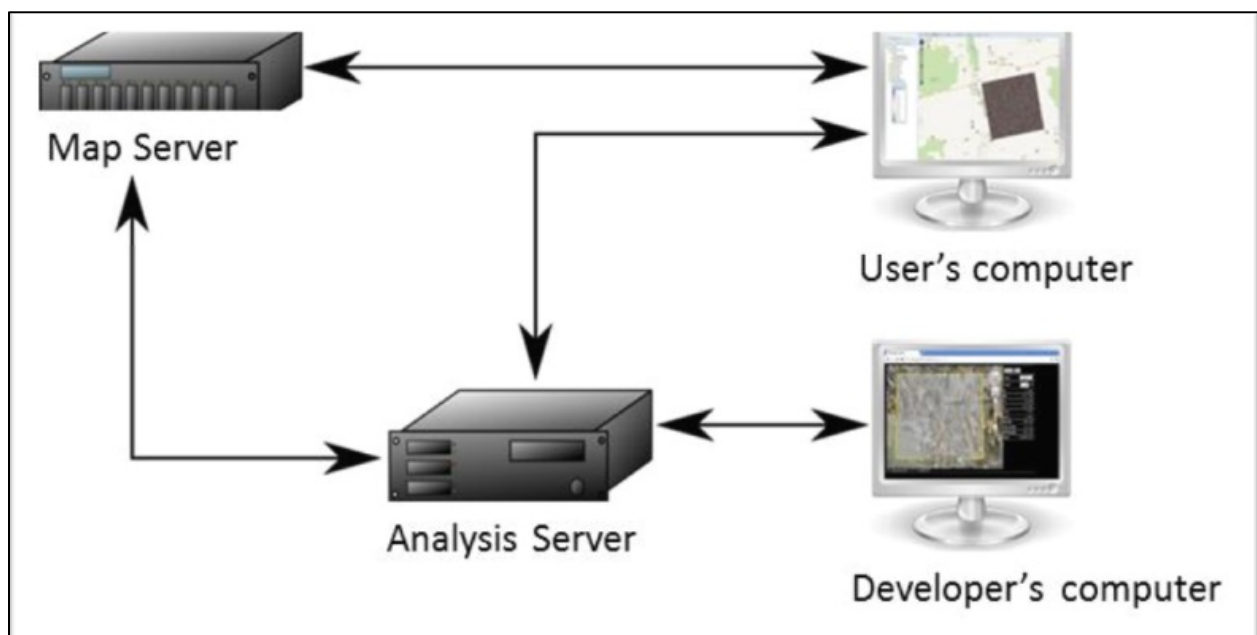


Figure 6. Simplified system diagram

4.2 Site Selection Algorithm

After the user makes selections for the weighting regime and presses the “run” button, the site selection program begins. The site selection program first generates new cost surface layers based on the weighting criteria the user selects. These cost surfaces depend on the underlying function in the weighting regime and fall into four general categories:

- Resource: Exactly one layer representing the energy resource in areal units (e.g., watts per square meter), as seen in Figure 7
- Areal Cost Contribution: One or more layers representing “per unit area” cost contributions (e.g., land acquisition cost) in unspecified areal units (e.g., dollars per acre), as seen in Figure 8

- **Lineal Cost Contribution:** Zero or more layers representing distance-related costs (e.g., distance from transmission lines) in unspecified “lineal” units (e.g., dollars per mile)
- **Mask:** One or more binary layers representing total exclusion zones (e.g., lakes, critical habitat).

The program then combines cost surfaces for individual layers into two cumulative cost surfaces representing areal and distance-based costs. The cumulative areal cost surface represents the combination of areal costs and masking layers, while the cumulative lineal cost surface represents the combination of lineal cost surfaces and masking layers. This cumulative areal cost surface is then “flooded” by iteratively applying incrementally higher cost values and removing those pixels whose values fall above it, leaving contiguous areas of minimum cost. Areas of land too small for the footprint of the facility, determined from the user inputs, including weighting criteria, are removed using a segmentation algorithm, leaving only areas that meet all criteria (Figure 9). If no areas are large enough to accommodate the facility, the threshold level is increased and the process is repeated until candidate sites are found.

Finally, the map of possible locations based on areal costs, can be ranked using the lineal cost surface and a number of “best” solutions are sent back to the user (Figure 10).




Figure 7. Example solar resource data. Darker pixels indicate areas of higher quality solar resource, which can be defined using a variety of metrics (e.g., watts per square meter).




Figure 8. Example areal surface cost layer. Darker areas indicate higher costs. Multiple areal surface layers can be incorporated in the decision-making process to represent costs per unit of area (e.g., dollars per acre) for items such as land acquisition costs and site preparation costs.




Figure 9. Contiguous areas in included region. Areas shown are those that have successfully met all criteria (e.g., solar resource, land slope, transmission lines, existing solar energy development zones, roads, BLM Areas of Critical Environmental Concern (ACEC), critical plant habitat, and critical fauna habitat).



Figure 10. Top 10 segments contiguous locations by lineal cost for facility siting based on segments shown in Figure 9. The segments can be ranked according to costs and a subset can be displayed.

5 Discussion

The initial goals of this project were to carefully consider and build a system that could support the decision-making process for siting renewable energy development, particularly large-scale solar development. The project succeeded in conceptualizing elements critical to a decision support tool and developing a proof-of-concept.

5.1 Proof-of-Concept Tool

The generic nature and broad scope of the tool makes it suitable for use by a variety of user groups: policymakers, utility companies, developers of large-scale renewable energy projects, and NGOs. User-defined weights and preferences for individual layers allow these various groups to test and contrast and compare individual factors, or subsets of factors, that impact a siting decision or optimization question. In addition, utilization of the same tool by different stakeholders facilitates communication and allows direct comparisons of stakeholder priorities and the development of potential compromises. This tool can also be used to evaluate existing SEZs developed by other organizations. In addition, the tool can aid in the evaluation of expected cumulative environmental impacts of development, given different levels of solar project deployment. In some cases solar development could conflict with other types of energy or mineral development (e.g., areas with high wind or geothermal potential, uranium deposits); this tool provides the opportunity to weigh and manage such risks. Importantly, this tool provides a common metric by which different stakeholders can communicate and express their views in a tangible, quantitative manner.

5.2 Future Directions and Recommendations

Developing the proof-of-concept resulted in a number of suggestions for finalizing the tool, improving its performance, and enhancing its capabilities. These suggestions can be implemented relatively quickly and are not hurdles to finalizing the tool.

Currently, the two main components of the tool work well: The optimization software performs its functions and produces appropriate results while the software running the Web interface creates platform for user experience. The main hurdle was linking these two pieces of the software so that they ran on the same server and shared inputs and outputs. Implementing the analysis service in the proof-of-concept phase proved more difficult than originally envisioned. We underestimated the large number of dependencies on external libraries and the challenge of integrating them smoothly. A complete implementation of the tool will use more external libraries. The service should preferably support the Open Geospatial Consortium (OGC) Web Coverage Service (WCS) rather than Web Map Service (WMS), and the Geospatial Data Abstraction Library (GDAL) for GeoTIFF, and should directly support a geo-enabled relational database management system (RDBMS), such as PostGIS. These were omitted from the proof-of-concept phase due to time constraints.

To link the two pieces of software and increase the tool's performance, we recommend that future versions of this tool use modified software architecture. Computation on layers should be adaptive and based on the historical pattern of visitor demands in a Web-based platform, which might better facilitate non-linear layer combinations and cost calculation. It should also be possible to specify a formula for layer cost calculation and combination in a configuration file

written in a general language, such as MathML, which could then be just-in-time (JIT) compiled into fast object code. The possibility of implementing the service as an extension to an existing open source OGC WMS/WCS server, such as GeoServer, should also be investigated.

In addition to software optimization, we recommend increasing the number of variables used in the optimization equation. This would improve the siting assessment functionality. For example, data regarding water availability, existing land use, land ownership, local permitting requirements and regulations, visual impact concern, and cultural impact potential could be incorporated into the tool. Incorporating such factors would improve the ability of this tool to inform real-life siting decisions.

Because the tool is designed to facilitate discussion among stakeholders, we need feedback from users regarding its use and functionality. Understanding how this prototype tool elicits feedback from multiple stakeholder groups, such as policymakers, public utility commission representatives, solar developers, and environmental groups would greatly improve the scope and directions of future versions. Different stakeholder groups are likely to have different recommendations for improvement, which could be implemented in the next version of this tool.

Solar development would likely face fewer barriers to deployment if sites were selected on non-productive or already disturbed land parcels. These lands could be better defined and characterized through improvements in remote sensing techniques and further research. Such data could be incorporated into the tool and into the optimization equation.

This tool could be expanded to address development of other fuels and technologies, such as wind, geothermal, biomass, natural gas, oil, and uranium. Detailed datasets could be incorporated into the tool to analyze optimum siting for these energy developments. In addition, the tool could also allow for comparisons between different energy technologies or identify co-location opportunities for various technologies.

6 Conclusions

We developed a proof-of-concept solar siting tool that goes beyond previous siting criteria analyses by solving a transparent, user-definable, optimization algorithm that is Web-based and interactive. This project lays the foundation for future work and analyses related to determining suitable locations for new solar energy developments. The tool, if developed further, could assist policymakers and industry representatives in their solar project siting decisions by giving all stakeholders a common framework by which they can discuss economic and environmental tradeoffs associated with different locations. Furthermore, the tool could be expanded to include other energy and mineral resource development opportunities, such that multiple resource development options could be evaluated at once and compared in terms of their economic, environmental, and social impacts.

References

- BLM/DOE. (2010). *Draft Solar Peis: Appendix M: Methodologies and Data Sources for the Analysis of Impacts of Solar Energy Development on Resources*. DOE/EIS-0403. Washington, DC: U.S. Department of Energy.
- BLM/DOE. (2012). *Final Programmatic Environmental Impact Statement (PEIS) for Solar Energy Development in Six Southwestern States*. Report No: FES 12-14; DOE/EIS-0403. Washington, DC: U.S. Department of the Interior.
- CEC. (2009). *Best Management Practices and Guidance Manual: Desert Renewable Energy Projects*. CEC-700-2009-016-SDREV. Sacramento, CA: California Energy Commission.
- Clark, T.W. (2002). *The Policy Process: A Practical Guide for Natural Resource Professionals*. New Haven, CT: Yale University Press.
- DOE. (2012). *SunShot Vision Study*. Report No.: DOE/GO-102012-3037. Washington, D.C.: Office of Energy Efficiency and Renewable Energy, U.S. Department of Energy.
- MacDonald, G.M. (2010). "Water, Climate Change, and Sustainability in the Southwest." *Proceedings of the National Academy of Sciences* (107:50); pp. 21256-21262.
- Mehos, M.; Kabel, D.; Smithers, P. (2009). "Planting the Seed." *IEEE Power and Energy Magazine* (7:3); p. 7.
- SEIA. (2014). *Major Solar Projects in the United States Operating, Under Construction, or Under Development*. Washington, DC: Solar Energy Industries Association. Updated March 6, 2014.
- TWC. (2009). *Comments to California Energy Commission: Siting Criteria for Re Pilot Projects*. Oak Glen, CA: The Wildlands Conservancy.
- TWS. (2009). *Comments to the Solar Energy Peis-Solar Energy Study Areas: Scoping Comments on the Solar Energy Study Areas for the Solar Peis*. Oak Glen, CA: The Wilderness Society.
- WGA/DOE. (2009). *Western Renewable Energy Zones - Phase 1 Report*. Washington, DC: U.S. Department of Energy.

Appendix: System Technical Description

System Structure

Data were loaded into a spatially enabled PostgreSQL database. Once loaded into the database, the data were publicly exposed as OGC-compliant Web Mapping Services (WMS) using GeoServer for use by the Web Analysis Service (WAS).

Figure A-1 shows a simplified Unified Modeling Language (UML) class diagram of the web analysis service; certain classes, such as exceptions, have been omitted for clarity. Essential supporting library packages, which are largely opaque, are shown in yellow, and “Request,” effectively the entry point, not implemented as a class per se, is shown in blue.

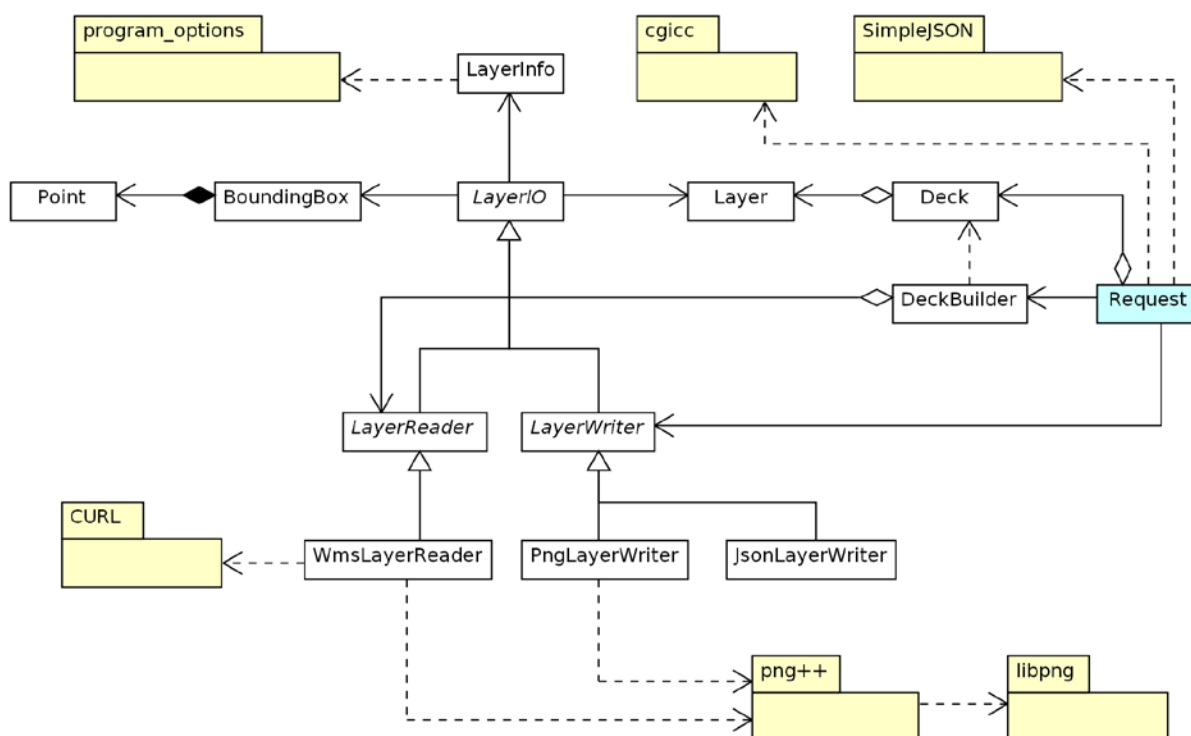


Figure A-1. Simplified Unified Modeling Language (UML) class diagram of tool

For every service “Request,” a “DeckBuilder” is instantiated, which builds a “Deck” of “Layers” using both information parsed from the request via “cgicc” and from a configuration file, via “boost::program_options.” Each Deck contains a number of distinct types of layers (not represented in Figure A1, including resource, areal cost contribution, lineal cost contribution, mask, total area costs, and total lineal costs).

Although “WmsLayerReader” is not the only “LayerReader”—a “FileLayerReader” and a “RandomLayerReader” exist for debugging and demonstration purposes — it is the “LayerReader” that provides the crucial analytical capability. It is sufficiently general that, in

principle, other concrete layer readers (e.g. one to read *GeoTIFF* files via *GDAL*) could be implemented and integrated into the service.

Similarly, the “PngLayerWriter” and “JsonLayerWriter” are the current operating “LayerWriters,” but “LayerWriters” for other output formats could be easily implemented as wrappers around existing application programming interfaces (API) supporting those formats; the major obstacle here would be the complexity of the chosen API. The WAS implementation currently stands at about 2,200 non-comment lines of C++.

Operation

An incoming service request is parsed for client-supplied input parameters such as geographic bounding box, spatial reference system (SRS) identifier, and resolution. The raw string forms of these parameters are hashed using the MD5 cryptographic message digest algorithm to create a 128-bit identifier with a very high probability of being unique. This identifier is a cache key: if an identical request has previously been serviced, the response is served from an output file cache and no further processing takes place.

If there is no cache entry for the request, a “DeckBuilder” object is initiated, given the input parameters and the name of a configuration file. The configuration file is parsed for information about each layer, including type, weight, URL of the WMS server, etc. This information is augmented by the input parameters (SRS, bounding box, resolution) to form a WMS request for each layer. The WMS request is hashed, again using the MD5 algorithm, and the request is serviced from the input cache if possible. Otherwise, a WMS request is made to obtain raw layer data.

Because the bounding box and cell-size resolution are identical for each raster data layer, all of the layers are vertically integrated geospatially. Areal and lineal cost contribution layers are linearly combined according to weights specified in the configuration file (or specified as input parameters) to yield the two total cost layers. The maximum and minimum values of the two layers are recorded.

The areal total cost layer is then linearly discretized to 250 integer values. Mask layers are then applied and excluded pixels values are set to 255. Values between 251 and 254 inclusive are not used, and should never occur. This surface is then flooded by comparing the values to a threshold, increasing from 0. A binary solution layer records those pixels that are above (0x00) and below (0xff) the threshold. A de-speckling algorithm removes isolated pixels, and a segmentation algorithm identifies contiguous regions within the non-excluded region.

The available resource for each region is evaluated (e.g., the summation of the product of area and energy density) and those regions too small to satisfy the resource requirement are eliminated. If no region can satisfy the resource requirement, the flood level is incremented and a new solution is sought until one or more regions satisfy the basic resource requirement. The solution is then ranked by minimum lineal cost in the region, and “best” candidate facility siting solutions are presented.

Google Earth Development Client

For development purposes, an AJAX-based client that uses the Google Earth browser plugin was written; it currently stands at approximately 1,000 non-comment lines of JavaScript, HTML (HyperText Markup Language), and CSS (Cascading Style Sheets). In outline, it enables a curvilinear square bounding box to be selected in the Google Earth plugin, and by clicking “Go,” submits a request to the analysis service, which returns a .PNG image, which is then added to Google Earth as an overlay. Currently, the resolution of this image is fixed. Experiments with checkerboard images indicate that, were it to be scaled to the resolution actually required, it would have to be generalized by approximately a factor of 4 before being passed to Google Earth owing to scaling artifacts. In other words, although the calculations could (and, indeed, should) be performed at the needed resolution, the results should ideally be generalized or aggregated in a post-processing stage. The Google Earth plugin API requires the selection region to be drawn as a sequence of points representing a closed polygon; then, on every mouse-click event, the polygon must be deleted and redrawn anew. Thus, the responsiveness of this client is not 100% smooth, and remains somewhat “clunky” due to the challenge of implementing dynamic box selection functionality.

OpenCarto User Interface

A web-mapping application was developed to provide a publicly accessible and functionally robust interface to facilitate user interaction with the WAS, as well as to visually display analysis results. This user interface was built upon the existing NREL-developed web-mapping framework called OpenCarto. OpenCarto is a collection of both open-source and custom software tools that provide data management, querying, visualization, and analysis capabilities for spatial data in a common template-based web application. In short, OpenCarto provides a set of tools to easily display and interact with spatial data on the web.

OpenCarto also provides a highly extensible ExtJS-based platform that makes it easy to create application-specific user interface elements. This aspect of OpenCarto was leveraged to create a custom window to capture user-defined cost calculation layer preferences and package these inputs in a web request to the WAS.